arXiv:1603.07779v1 [cs.CC] 24 Mar 2016

Nonnegative Rank vs. Binary Rank

Thomas Watson*

March 28, 2016

Abstract

Motivated by (and using tools from) communication complexity, we investigate the relationship between the following two ranks of a 0-1 matrix: its nonnegative rank and its binary rank (the log of the latter being the unambiguous nondeterministic communication complexity). We prove that for partial 0-1 matrices, there can be an exponential separation. For total 0-1 matrices, we show that if the nonnegative rank is at most 3 then the two ranks are equal, and we show a separation by exhibiting a matrix with nonnegative rank 4 and binary rank 5, as well as a family of matrices for which the binary rank is 4/3 times the nonnegative rank.

1 Introduction

Let M be a total 0-1 matrix of size $n \times m$. We consider three notions of the rank of M (over the reals).

(i) The familiar *real rank*, denoted $\text{rank}(M)$, is the smallest r for which there exists a *product* decomposition $M = UV$ where U is $n \times r$ and V is $r \times m$. Equivalently, it is the smallest r for which there exists a *sum* decomposition $M = M^{(1)} + \cdots + M^{(r)}$ where each component is rank 1, say $M^{(i)} = u^{(i)}v^{(i)}$ where $u^{(i)}$ is $n \times 1$ and $v^{(i)}$ is $1 \times m$.

(ii) The *nonnegative rank*, denoted $\text{rank}_+(M)$, is defined in the same way but with the restriction that U and V are entry-wise nonnegative, or equivalently that all $u^{(i)}, v^{(i)}$ (and hence $M^{(i)}$) are entry-wise nonnegative.

(iii) The *binary rank*, denoted $\text{rank}_{0,1}(M)$, is defined in the same way but with the further restriction that all entries of U and V come from $\{0,1\}$, or equivalently that all entries of $u^{(i)}, v^{(i)}$ (and hence $M^{(i)}$) come from $\{0,1\}$.

We define a *partial* 0-1 matrix as having entries from $\{0,1,*\}$, where $*$ entries are wildcards representing arbitrary numbers (from some set depending on the context; see details below). The real, nonnegative, or binary rank of a partial 0-1 matrix M are defined (respectively) as the minimum real, nonnegative, or binary rank of any total real matrix that agrees with M on the non-$*$ entries. In the case of nonnegative rank, the total matrix must have nonnegative entries, and in the case of binary rank, the total matrix must have nonnegative integer entries (but in any case, its entries corresponding to $*$ need not come from $\{0,1\}$). Note that for all 0-1 matrices M (total or partial), $\text{rank}(M) \leq \text{rank}_+(M) \leq \text{rank}_{0,1}(M)$.

*Department of Computer Science, University of Toronto. Supported by funding from NSERC.

These ranks are closely related to measures of communication complexity of the two-party function associated with M. (See [KN97, Juk12] for background on communication complexity.) The binary rank is just the number of (combinatorial) rectangles needed to partition the 1's of M. This corresponds to unambiguous nondeterminism (the communication complexity analogue of the classical complexity class UP), so the unambiguous nondeterministic communication complexity[1] is defined to be $\log_2 \operatorname{rank}_{0,1}(M)$.[2]

Similarly, $\log_2 \operatorname{rank}(M)$ can be viewed as a communication measure corresponding to the classical complexity class LWPP (representing randomized algorithms whose acceptance probability is one common value on all 0-inputs and a different common value on all 1-inputs). The famous *Log-Rank Conjecture* asserts that $\log \operatorname{rank}(M)$ is polynomially related to the deterministic communication complexity of M, for every total 0-1 matrix M.

Analogously to the correspondences between $\log_2 \operatorname{rank}_{0,1}$ and UP, and between $\log_2 \operatorname{rank}$ and LWPP, we can view $\log_2 \operatorname{rank}_+$ as corresponding to a classical complexity class representing randomized algorithms whose acceptance probability is 0 on all 0-inputs and a positive common value on all 1-inputs (though this class appears to lack a name). Nonnegative rank of a total 0-1 matrix M also corresponds to the following random sampling problem: Alice and Bob are given private randomness but no input, and their goal is for Alice to output a row index and Bob to output a column index such that the corresponding entry of M is uniformly distributed over all the 1's of M.[3] It turns out that one-way communication is optimal in this setting, and the nonnegative rank is exactly the minimum number of transcripts needed to solve this problem (see, e.g., [JSWZ13], where it is called "correlation complexity"), and hence the minimum number of bits of communication is $\log_2 \operatorname{rank}_+(M)$. If M is a partial 0-1 matrix, then the associated sampling problem is for Alice's row and Bob's column to be uniformly distributed over the 1 entries of M, *conditioned* on being a non-$*$ entry.

We consider the relationship between nonnegative rank and binary rank (hence between the above sampling problem and unambiguous nondeterminism). In terms of motivation, nonnegative rank of real matrices has myriad applications in theoretical computer science, other branches of computer science, and other scientific disciplines (see [CR93, Gil14]), but there seems to be a dearth of proof techniques that directly exploit special properties of nonnegative rank. For example, there seem to be almost no known general-purpose techniques for upper bounding rank_+ without upper bounding $\operatorname{rank}_{0,1}$, or for lower bounding $\operatorname{rank}_{0,1}$ without lower bounding rank_+. We provide some (albeit ad hoc) techniques that differentiate between rank_+ decompositions and $\operatorname{rank}_{0,1}$ decompositions. Finally, we feel it is always natural to compare different measures of complexity, and from a purely combinatorial point of view it is natural to compare different notions of rank of 0-1 matrices.

[1] This is closely related to the clique vs. independent set family of problems, since it is well-known that every two-party total function reduces to the clique vs. independent set problem for a graph on $\operatorname{rank}_{0,1}(M)$ nodes.

[2] There are two other types of rank defined in the same way as binary rank but with different arithmetic. With boolean arithmetic (so + is OR) this is sometimes known as *boolean rank*; it is the number of rectangles needed to cover the 1's of M; it lower bounds the nonnegative rank and corresponds to NP. With $GF(2)$ arithmetic (so + is XOR) this is sometimes known as *xor rank*; it is the number of rectangles needed to cover each 1 of M an odd number of times and each 0 of M an even number of times; it lower bounds the real rank and corresponds to $\oplus$P.

[3] A closely related type of problem (which was studied in the communication complexity setting in [ASTS+03], and in the setting of constant-depth circuits in [Vio12]) is to sample a uniformly random input-output pair of a given function. In our problem we just ask to sample a uniformly random 1-input, which can be viewed as peripherally motivated by the topic of uniform sampling / approximate counting of NP witnesses.

It is known that for every total 0-1 matrix M,

$$\begin{aligned}\log_2 \operatorname{rank}_{0,1}(M) &\leq \text{deterministic communication complexity of } M \\ &\leq O\big((\log \operatorname{rank}(M)) \cdot (\text{nondeterministic communication complexity of } M)\big) \\ &\leq O(\log^2 \operatorname{rank}_+(M)) \\ &\leq O(\log^2 \operatorname{rank}_{0,1}(M)) \qquad (1)\end{aligned}$$

where the second line follows by a result of [LS93]. (A simple proof of the special case "deterministic communication complexity of $M \leq O(\log^2 \operatorname{rank}_{0,1}(M))$" was given earlier in [Yan91].) We show that the quadratic upper bound $\log \operatorname{rank}_{0,1}(M) \leq O(\log^2 \operatorname{rank}_+(M))$ does not hold in general for partial 0-1 matrices; there is an exponential separation. Our proof of this exploits machinery from [GLM+15, Göö15] in a new way and also (perhaps surprisingly) exploits the upper bound (1) for total 0-1 matrices. Prior to this work it was open whether nonnegative rank and binary rank are *always equal* for total 0-1 matrices; we refute this, although our counterexamples are very far from resolving whether the quadratic upper bound is tight.

Theorem 1. *There exists a family of partial* 0-1 *matrices* M *such that* $\operatorname{rank}_{0,1}(M) \geq 2^{\operatorname{rank}_+(M)^{\Omega(1)}}$.

Theorem 2.

(1) For every total 0-1 *matrix* M*, if* $\operatorname{rank}_+(M) \leq 3$ *then* $\operatorname{rank}_+(M) = \operatorname{rank}_{0,1}(M)$.

(2) There exists a total 0-1 *matrix* M *such that* $\operatorname{rank}_+(M) = 4$ *and* $\operatorname{rank}_{0,1}(M) = 5$.

(3) There exists a total 0-1 *matrix* M *such that* $\operatorname{rank}_+(M) = 9$ *and* $\operatorname{rank}_{0,1}(M) = 12$.

Tensoring the matrix from Theorem 2(3) with the identity shows that for every k divisible by 9 there exists a total 0-1 matrix with nonnegative rank k and binary rank $(4/3)k$. After taking logs, this yields an extremely meager gap between the sampling and unambiguous nondeterminism complexity measures (an additive constant less than 1, which would sometimes be wiped out by taking ceilings, anyway). The primary open question is whether a better gap can be exhibited. A natural approach is to amplify the gap by repeatedly tensoring a matrix with itself. Lower bounding the binary rank of the resulting matrix seems related to the notoriously difficult direct sum problem for deterministic communication complexity (since the k-fold tensor product corresponds to taking the AND of the outputs of k independent instances of the original two-party function, which is no harder than computing the values of all k outputs). We have been unable to prove any interesting such lower bounds, even for the concrete matrices from Theorem 2(2) and Theorem 2(3). This is further discussed in Section 4.

There are several other works investigating the relationships of different ranks with each other and with measures of communication complexity (for total matrices). Göös, Pitassi, and Watson [GPW15] exhibited a nearly quadratic (essentially tight) gap between deterministic communication complexity and the log of the binary rank, as well as a nearly power 1.5 gap between deterministic communication complexity and the log of the number of monochromatic rectangles needed to partition the 0's and 1's of the matrix (i.e., the log of the sum of the binary rank and the binary rank of the complement matrix). That power 1.5 gap (which has subsequently been improved to quadratic by Ambainis, Kokainis, and Kothari [AKK16]) improves the factor 2 gap due to Kushilevitz, Linial, and Ostrovsky [KLO99], which was also the previous record for deterministic communication complexity vs. log of the binary rank (and we mention that for the latter,

Kushilevitz and Weinreb [KW09] had used different techniques to obtain a weaker factor 1.12 gap). For deterministic communication complexity vs. log of the real rank, the best gap before [GPW15] was a power $\log_3 6 \approx 1.63$ gap due to Nisan, Wigderson, and Kushilevitz [NW95]; in the other direction, Lovett [Lov14] showed that the deterministic communication complexity of any M is $O\big(\sqrt{\operatorname{rank}(M)}\log\operatorname{rank}(M)\big)$. Göös [Göö15] exhibited matrices for which the conondeterministic communication complexity is at least a power 1.128 greater than the log of the binary rank, improving on the factor 2 gap due to Shigeta and Amano [SA15]. Shitov [Shi14] showed that for every $n \times m$ real nonnegative matrix with real rank 3, the nonnegative rank is at most $\lceil \frac{6}{7} \cdot \min(n,m) \rceil$.

2 Partial Matrices

In this section we prove Theorem 1. Letting $g\colon \{0,1\}^b \times \{0,1\}^b \to \{0,1\}$ be a total function (usually called a *gadget*), define $g^n\colon \{0,1\}^{bn} \times \{0,1\}^{bn} \to \{0,1\}^n$ by $g^n(x,y) = \big(g(x_1,y_1),\dots,g(x_n,y_n)\big)$ where $x = (x_1,\dots,x_n)$ and $y = (y_1,\dots,y_n)$ with $x_i, y_i \in \{0,1\}^b$ for each i. For any (possibly partial) function $f\colon \{0,1\}^n \to \{0,1\}$, the (possibly partial) two-party communication problem $f \circ g^n\colon \{0,1\}^{bn} \times \{0,1\}^{bn} \to \{0,1\}$ is defined by $(f \circ g^n)(x,y) = f\big(g(x_1,y_1),\dots,g(x_n,y_n)\big)$, and we identify this with a $2^{bn} \times 2^{bn}$ partial 0-1 matrix with rows indexed by x and columns indexed by y.

Consider the partial function $f\colon \{0,1\}^n \to \{0,1\}$ (where n is even) defined by

$$f(z) \;=\; \begin{cases} 1 & \text{if } |z| = n/2 \\ 0 & \text{if } |z| = 0 \end{cases} \tag{2}$$

where $|z|$ denotes the Hamming weight of z.

Lemma 1. *For f as in (2) and for all $g\colon \{0,1\}^b \times \{0,1\}^b \to \{0,1\}$, we have $\operatorname{rank}_+(f \circ g^n) \le n \cdot 2^b$.*

Proof. For $i \in \{1,\dots,n\}$ and $s \in \{0,1\}^b$, define the total function (matrix) $M^{(i,s)}\colon \{0,1\}^{bn} \times \{0,1\}^{bn} \to \mathbb{R}_{\ge 0}$ by

$$M^{(i,s)}(x,y) \;=\; \begin{cases} 2/n & \text{if } x_i = s \text{ and } g(x_i,y_i) = 1 \\ 0 & \text{otherwise} \end{cases}.$$

Since each $M^{(i,s)}$ is nonnegative and has rank 1, it suffices to check that $\sum_{i,s} M^{(i,s)}$ agrees with $f \circ g^n$ on the latter function's domain. If (x,y) is such that $|g^n(x,y)| = 0$ then $M^{(i,s)}(x,y) = 0$ for all (i,s). If (x,y) is such that $|g^n(x,y)| = n/2$ then $M^{(i,s)}(x,y) = 2/n$ for exactly $n/2$ many (i,s); specifically, there are $n/2$ many i's for which $g(x_i,y_i) = 1$, and for each such i there is exactly one s for which $x_i = s$. □

The paper [GLM+15] introduced and studied what we call the *confounding* gadget $g\colon \{0,1\}^b \times \{0,1\}^b \to \{0,1\}$, defined by

$$g(x_i,y_i) \;=\; \langle x_i, y_i \rangle \bmod 2 \qquad \text{and} \qquad b \;=\; b(n) \;=\; 100\log_2 n. \tag{3}$$

Lemma 2. *For f as in (2) and g as in (3), we have $\operatorname{rank}_{0,1}(f \circ g^n) \ge 2^{\Omega(\sqrt{n \log n})}$.*

Theorem 1 follows immediately from Lemma 1 (which shows that $\operatorname{rank}_+(f \circ g^n) \le n^{101}$ for the confounding gadget g) and Lemma 2. We conjecture that when g is the AND gadget (with $b = 1$), $\operatorname{rank}_{0,1}(f \circ g^n) \ge 2^{\Omega(n)}$, but we are unable to prove this.

To prove Lemma 2, we use the following property of the confounding gadget, which is a special case of the main technical result in [GLM+15] (and a streamlined, self-contained proof of this special case appears in [Göö15]). A *subcube* of codimension d is defined as the subset of $\{0,1\}^n$ consisting of all 2^{n-d} strings consistent with some partial assignment that fixes d of the bit positions. For $z \in \{0,1\}^n$, we also define the set $W_z \subseteq \{0,1\}^{bn} \times \{0,1\}^{bn}$ to be $(g^n)^{-1}(z)$.

Lemma 3. *For g as in (3), for every rectangle $X \times Y \subseteq \{0,1\}^{bn} \times \{0,1\}^{bn}$ and every $z \in \{0,1\}^n$, if $|(X \times Y) \cap W_z| \geq 2^{-c} \cdot |W_z|$, then there exists a subcube $Z \subseteq \{0,1\}^n$ of codimension $O(c/b) = O(c/\log n)$ such that $Z \subseteq g^n(X,Y)$ (i.e., for every $z' \in Z$ there exists an $(x,y) \in X \times Y$ such that $g^n(x,y) = z'$).*

Proof of Lemma 2. Suppose for contradiction there is a collection of rectangles $R_i = X_i \times Y_i$ for $i \in \{1, \ldots, 2^{o(\sqrt{n \log n})}\}$ such that if $(f \circ g^n)(x,y) = 1$ then $(x,y) \in R_i$ for exactly one i, and if $(f \circ g^n)(x,y) = 0$ then $(x,y) \in R_i$ for no i's.

First we claim that for all $i < j$, $|R_i \cap R_j| \leq 2^{-\Omega(n \log n)} \cdot 2^{2bn}$. Supposing not, we would have $|R_i \cap R_j \cap W_z| \geq 2^{-o(n \log n)} \cdot |W_z|$ for some z (since the sets W_z partition $\{0,1\}^{bn} \times \{0,1\}^{bn}$). Then by Lemma 3 there would exist a subcube $Z \subseteq g^n(R_i \cap R_j)$ of codimension $o((n \log n)/\log n) = o(n) \leq n/2$. But since any subcube of codimension $\leq n/2$ contains a 1-input of our f, this means $R_i \cap R_j$ contains a 1-input of $f \circ g^n$, contradicting our assumption. This proves the claim.

In particular, for all $i < j$ we have either $|X_i \cap X_j| \leq 2^{-\Omega(n \log n)} \cdot 2^{bn}$ or $|Y_i \cap Y_j| \leq 2^{-\Omega(n \log n)} \cdot 2^{bn}$; define $Q_{i,j} = (X_i \cap X_j) \times \{0,1\}^{bn}$ in the former case and $Q_{i,j} = \{0,1\}^{bn} \times (Y_i \cap Y_j)$ in the latter case. Then we have $|Q_{i,j}| \leq 2^{-\Omega(n \log n)} \cdot 2^{2bn}$ and hence $\left|\bigcup_{i<j} Q_{i,j}\right| \leq \sum_{i<j} |Q_{i,j}| \leq (2^{o(\sqrt{n \log n})})^2 \cdot 2^{-\Omega(n \log n)} \cdot 2^{2bn} \leq 2^{-\Omega(n \log n)} \cdot 2^{2bn}$.

Define $S = \left(\{0,1\}^{bn} \times \{0,1\}^{bn}\right) \setminus \bigcup_{i<j} Q_{i,j}$, and note that S is a rectangle. Define the total function $F \colon S \to \{0,1\}$ by $F(x,y) = 1$ iff $(x,y) \in R_i$ for some i; note that F is consistent with $f \circ g^n$. Since the rectangles $R_i \cap S$ are pairwise disjoint, we have $\log \mathrm{rank}_{0,1}(F) \leq o(\sqrt{n \log n})$ and thus the deterministic (in particular, conondeterministic) communication complexity of F is $o(n \log n)$ by (1).[4] Thus the set $W_{0^n} \cap S$ can be covered with $2^{o(n \log n)}$ many subrectangles of S that are disjoint from all R_i's. At least one of these subrectangles, call it T, covers at least a $2^{-o(n \log n)}$ fraction of $W_{0^n} \cap S$. Since $|W_{0^n}| \geq 2^{-n} \cdot 2^{2bn}$, we have $|W_{0^n} \cap S| \geq |W_{0^n}| - \left|\bigcup_{i<j} Q_{i,j}\right| \geq \frac{1}{2} \cdot |W_{0^n}|$. Thus T also covers at least a $2^{-o(n \log n)}$ fraction of W_{0^n}. By Lemma 3, there exists a subcube $Z \subseteq g^n(T)$ of codimension $o((n \log n)/\log n) = o(n) \leq n/2$. But since any subcube of codimension $\leq n/2$ contains a 1-input of our f, this means T contains a 1-input of $f \circ g^n$, contradicting the fact that T is disjoint from all R_i's. □

3 Total Matrices

In this section we prove Theorem 2. All matrices are tacitly assumed to be total in this section.

3.1 A Lemma

We use the following lemma in the proof of Theorem 2(1) (but we remark that this is not the only property of nonnegative rank we use in the proof of Theorem 2(1)).

[4] Any improvement in the cost of converting unambiguous protocols to conondeterministic protocols (for total functions) would yield a corresponding quantitative improvement in Lemma 2; however, the result of [Göö15] shows that the exponent of 2 in such a conversion cannot be decreased below 1.128.

The *support* of a matrix is the set of locations of nonzero entries. The support of a rank 1 matrix is always a rectangle.

Lemma 4. *For every* 0-1 *matrix* M*, every* rank_+ *decomposition* $M = M^{(1)} + \cdots + M^{(r)}$ *with supports* $R_1, \ldots, R_r$*, and every* $i \in \{1, \ldots, r\}$*,* $R_i \setminus \bigcup_{j \neq i} R_j$ *is a rectangle.*

Proof. Denote $R = R_i \setminus \bigcup_{j \neq i} R_j$. Assume $(a, b), (c, d) \in R$ and $a \neq c$ and $b \neq d$. Then we have $M^{(i)}_{a,b} = M^{(i)}_{c,d} = 1$, and since $M^{(i)}$ has rank 1 we have $M^{(i)}_{a,d} \cdot M^{(i)}_{c,b} = M^{(i)}_{a,b} \cdot M^{(i)}_{c,d} = 1$. Since every entry of $M^{(i)}$ is in the range $[0, 1]$, this forces $M^{(i)}_{a,d} = M^{(i)}_{c,b} = 1$ and so $(a, d), (c, b) \in R_i$. In turn, this forces $M^{(j)}_{a,d} = M^{(j)}_{c,b} = 0$ for all $j \neq i$ and hence $(a, d), (c, b) \notin R_j$. Thus $(a, d), (c, b) \in R$ and so R is a rectangle. □

As an aside (not needed for Theorem 2(1)), here is a simple application illustrating Lemma 4.

Observation 1. *Define the intersection graph of a collection of rectangles to have nodes representing the rectangles, and an edge between two nodes iff their rectangles intersect. Consider a* 0-1 *matrix* M *with* $\text{rank}_+(M) = r$ *and a* rank_+ *decomposition* $M = M^{(1)} + \cdots + M^{(r)}$ *with supports* $R_1, \ldots, R_r$*. If the intersection graph of the supports is bipartite, then* $\text{rank}_+(M) = \text{rank}_{0,1}(M)$*.*

Proof. Consider a set of rectangles $\{R_i\}_{i \in S}$ that simultaneously forms an independent set and a vertex cover. For $i \notin S$, define $R'_i = R_i \setminus \bigcup_{j \in S} R_j$. We claim that the collection $\{R_i\}_{i \in S} \cup \{R'_i\}_{i \notin S}$ forms a partition of the 1's of M into r rectangles. *Coverage:* Since the collection of all R_i's forms a cover of the 1's of M, it follows by definition that $\{R_i\}_{i \in S} \cup \{R'_i\}_{i \notin S}$ also forms a cover. *Disjointness:* The rectangles R_i for $i \in S$ are disjoint from each other by the independent set property. The rectangles R_i for $i \notin S$ are disjoint from each other by the vertex cover property, so certainly the sets R'_i for $i \notin S$ are disjoint from each other. For $i \notin S$ and $j \in S$, R'_i and R_j are disjoint by definition. *Rectangles:* For $i \notin S$, we have $R'_i = R_i \setminus \bigcup_{j \neq i} R_j$ by the vertex cover property, and the latter set is a rectangle by Lemma 4. □

3.2 Proof of Theorem 2(1)

Consider the following two techniques for converting a 1's cover of M by rectangles $R_1, \ldots, R_r$ into a 1's partition of M.

Technique 1: For some permutation π on $\{1, \ldots, r\}$, use the sets $R_i \setminus \bigcup_{j : \pi(j) < \pi(i)} R_j$ (for $i \in \{1, \ldots, r\}$).

Technique 2: Let $R_i = A_i \times B_i$. For $\alpha \in \{1, -1\}^r$, we say that $\bigcap_i A_i^{\alpha_i}$ is a *type* of row (with respect to the cover $R_1, \ldots, R_r$), where $A_i^1 = A_i$ and $A_i^{-1} = \overline{A_i}$. We say the type is *nontrivial* if α is not all -1's. Note that all rows of the same type are identical, and thus each nontrivial nonempty type has an associated rectangle that covers the 1's in all rows of that type. The technique is to use this collection of rectangles, over all the nontrivial nonempty types of rows, or do the analogous thing for nontrivial nonempty types of columns.

Technique 1 is guaranteed to produce a small partition, but the sets are not guaranteed to be rectangles. Technique 2 is guaranteed to produce a partition into rectangles, but it is not guaranteed to be a small partition (it is small if many types of rows/columns are empty). Our approach to

prove Theorem 2(1) is to argue that, for a cover arising as the supports of a nonnegative rank 3 decomposition, *either* Technique 1 works *or* Technique 2 works.

In principle, Theorem 2(1) could be proved by brute force. The reason is because in a nonnegative rank 3 decomposition, there can be at most seven nontrivial types of rows and seven nontrivial types of columns, and duplicate rows and columns can be deleted without changing the nonnegative rank or the binary rank. Hence if there were a counterexample to Theorem 2(1), there would be a counterexample of size at most 7×7. However, such a brute force argument would be unenlightening; our argument provides insight into why Theorem 2(1) is true.

Definition 1. *We say a pair of rectangles (R, Q) is* compatible *if $Q \setminus R$ is a rectangle, in other words, either $R \cap Q = \emptyset$, or the rows of R are a superset of the rows of Q, or the columns of R are a superset of the columns of Q.*

Proof of Theorem 2(1). Of course, $\text{rank}_{0,1}(M) \geq \text{rank}_+(M)$ for all 0-1 matrices M, so we just need to show that if $\text{rank}_+(M) \leq 3$ then $\text{rank}_{0,1}(M) \leq \text{rank}_+(M)$. The case $\text{rank}_+(M) \leq 1$ is trivial. Suppose $\text{rank}_+(M) = 2$ and let $M = M^{(1)} + M^{(2)}$ be a rank_+ decomposition with supports R_1, R_2. By Lemma 4, $R_2 \setminus R_1$ is a rectangle, and hence R_1 and $R_2 \setminus R_1$ are two rectangles forming a 1's partition of M, so $\text{rank}_{0,1}(M) \leq 2$.

Suppose $\text{rank}_+(M) = 3$ and let $M = M^{(1)} + M^{(2)} + M^{(3)}$ be a rank_+ decomposition with supports R_1, R_2, R_3. First assume two of these rectangles form a compatible pair, say (R_1, R_2). Then $R_2 \setminus R_1$ is a rectangle, and $R_3 \setminus (R_1 \cup R_2)$ is a rectangle by Lemma 4. Hence R_1 and $R_2 \setminus R_1$ and $R_3 \setminus (R_1 \cup R_2)$ are three rectangles forming a 1's partition of M, so $\text{rank}_{0,1}(M) \leq 3$.

Now assume no two of the rectangles R_1, R_2, R_3 form a compatible pair. We show that either there are only three nontrivial nonempty types of rows, or there are only three nontrivial nonempty types of columns (which gives a 1's partition into three rectangles by Technique 2). Say $R_i = A_i \times B_i$. Since neither (R_1, R_2) nor (R_2, R_1) are compatible, the following sets are all nonempty: $A_1 \setminus A_2$, $A_1 \cap A_2$, $A_2 \setminus A_1$, $B_1 \setminus B_2$, $B_1 \cap B_2$, $B_2 \setminus B_1$. Now observe that either $(A_1 \cap A_2) \times (B_1 \setminus B_2) \subseteq R_3$ or $(A_1 \setminus A_2) \times (B_1 \cap B_2) \subseteq R_3$, since if not then letting $(a, b) \in \big((A_1 \cap A_2) \times (B_1 \setminus B_2)\big) \setminus R_3$ and $(c, d) \in \big((A_1 \setminus A_2) \times (B_1 \cap B_2)\big) \setminus R_3$, we have $(a, b), (c, d) \in R_1 \setminus (R_2 \cup R_3)$ but $(a, d) \in R_1 \cap R_2$ and hence $(a, d) \notin R_1 \setminus (R_2 \cup R_3)$, so $R_1 \setminus (R_2 \cup R_3)$ is not a rectangle, contradicting Lemma 4. Similarly, either $(A_1 \cap A_2) \times (B_2 \setminus B_1) \subseteq R_3$ or $(A_2 \setminus A_1) \times (B_1 \cap B_2) \subseteq R_3$.

Henceforth assume that $(A_1 \cap A_2) \times (B_1 \setminus B_2) \subseteq R_3$ (if $(A_1 \setminus A_2) \times (B_1 \cap B_2) \subseteq R_3$ then a symmetric argument applies). If $(A_2 \setminus A_1) \times (B_1 \cap B_2) \subseteq R_3$ then $B_1 \subseteq B_3$ and thus (R_3, R_1) would be compatible, so we may henceforth assume that $(A_1 \cap A_2) \times (B_2 \setminus B_1) \subseteq R_3$.

Consider three cases: $(A_1 \setminus A_2) \cap A_3 = \emptyset$ or $A_1 \setminus A_2 \subseteq A_3$ or neither. In the second case, $A_1 \subseteq A_3$ and thus (R_3, R_1) would be compatible. In the third case, we claim that we must have $B_1 \subseteq B_3$, and thus (R_3, R_1) would be compatible. To prove the claim, first note that $R_1 \setminus (R_2 \cup R_3) = \big((A_1 \setminus A_2) \times B_1\big) \setminus R_3$ since $(A_1 \cap A_2) \times (B_1 \setminus B_2) \subseteq R_3$. We have $\emptyset \subsetneq (A_1 \setminus A_2) \cap A_3 \subsetneq A_1 \setminus A_2$ (since we are assuming the third case) and, if we assume $B_1 \not\subseteq B_3$, we have $\emptyset \subsetneq B_1 \cap B_3 \subsetneq B_1$ (since $\emptyset \subsetneq B_1 \setminus B_2 \subseteq B_3$). Together, these imply that $\big((A_1 \setminus A_2) \times B_1\big) \setminus R_3$ is not a rectangle, contradicting Lemma 4. So, we may henceforth assume the first case, namely $(A_1 \setminus A_2) \cap A_3 = \emptyset$. Similarly, we may henceforth assume $(A_2 \setminus A_1) \cap A_3 = \emptyset$.

To finish the proof, we consider two cases:

(i) $B_1 \cap B_2 \cap B_3 \neq \emptyset$, or

(ii) $B_1 \cap B_2 \cap B_3 = \emptyset$.

First assume (i) holds. We claim that $B_1 \cap B_2 \subseteq B_3$ which, together with $B_1 \smallsetminus B_2 \subseteq B_3$, implies that $B_1 \subseteq B_3$ and thus (R_3, R_1) would be compatible (in fact, (R_3, R_2) would also be compatible). Let $b \in B_1 \cap B_2 \cap B_3$ and suppose for contradiction that some $d \in (B_1 \cap B_2) \smallsetminus B_3$. Let $a \in A_1 \smallsetminus A_2$ and $c \in A_1 \cap A_2$ and $e \in A_2 \smallsetminus A_1$. Since $(A_1 \smallsetminus A_2) \cap A_3 = \emptyset$ we have $(a,b), (a,d) \in R_1 \smallsetminus (R_2 \cup R_3)$ and thus $M^{(1)}_{a,b} = M^{(1)}_{a,d} = 1$. Since $M^{(1)}$ has rank 1, this implies that $M^{(1)}_{c,b} = M^{(1)}_{c,d}$. Similarly, $M^{(2)}_{c,b} = M^{(2)}_{c,d}$ (by using e in place of a). However, $(c,b) \in R_3$ (since $A_1 \cap A_2 \subseteq A_3$) and $(c,d) \notin R_3$ (since $d \notin B_3$), and thus $M^{(3)}_{c,b} > 0 = M^{(3)}_{c,d}$. Hence

$$M_{c,b} \;=\; M^{(1)}_{c,b} + M^{(2)}_{c,b} + M^{(3)}_{c,b} \;>\; M^{(1)}_{c,d} + M^{(2)}_{c,d} + M^{(3)}_{c,d} \;=\; M_{c,d}$$

which is a contradiction since $M_{c,b} = 1 = M_{c,d}$.

Now assume (ii) holds. In this case, $(R_1 \cup R_2) \cap R_3 = (A_1 \cap A_2) \times \big((B_1 \smallsetminus B_2) \cup (B_2 \smallsetminus B_1)\big)$. If $A_3 \smallsetminus (A_1 \cup A_2) \neq \emptyset$ and $B_3 \smallsetminus (B_1 \cup B_2) \neq \emptyset$ then $R_3 \smallsetminus (R_1 \cup R_2)$ is not a rectangle (since if $a \in A_3 \smallsetminus (A_1 \cup A_2)$, $b \in B_3 \smallsetminus (B_1 \cup B_2)$, $c \in A_1 \cap A_2$, and $d \in B_1 \smallsetminus B_2$, then $(a,d), (c,b) \in R_3 \smallsetminus (R_1 \cup R_2)$ but $(c,d) \in R_1$), contradicting Lemma 4. If $A_3 \smallsetminus (A_1 \cup A_2) = \emptyset$ then $A_3 = A_1 \cap A_2$ and thus there are only three nontrivial nonempty types of rows ($A_1 \cap \overline{A_2} \cap \overline{A_3}$ and $A_1 \cap A_2 \cap A_3$ and $\overline{A_1} \cap A_2 \cap \overline{A_3}$). On the other hand, if $B_3 \smallsetminus (B_1 \cup B_2) = \emptyset$ then $B_3 = (B_1 \smallsetminus B_2) \cup (B_2 \smallsetminus B_1)$ and thus there are only three nontrivial nonempty types of columns ($B_1 \cap \overline{B_2} \cap B_3$ and $B_1 \cap B_2 \cap \overline{B_3}$ and $\overline{B_1} \cap B_2 \cap B_3$). □

3.3 Some Examples

We now give some examples that elucidate features of the proof of Theorem 2(1).

$$\begin{bmatrix} 1 & 1 & 0 & 1 \\ 1 & 0 & 1 & 1 \\ 0 & 1 & 1 & 1 \\ 1 & 1 & 1 & 1 \end{bmatrix} \qquad \begin{bmatrix} 0 & 1 & 0 & 0 \\ 1 & 1 & 1 & 0 \\ 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 \end{bmatrix} \qquad \begin{bmatrix} 1 & 1 & 0 \\ 1 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 1 & 1 \end{bmatrix}$$

The first matrix above shows that Lemma 4 cannot possibly be the only property of rank_+ decompositions used to prove Theorem 2(1), because this matrix has a 1's cover of three rectangles satisfying the conclusion of Lemma 4 ($\{1,4\} \times \{1,2,4\}$ and $\{2,4\} \times \{1,3,4\}$ and $\{3,4\} \times \{2,3,4\}$), yet the binary rank is 4.

The second matrix above shows that Technique 2 is not sufficient to prove Theorem 2(1), because the binary rank is 3, yet there are four distinct nonzero rows and four distinct nonzero columns (hence necessarily four nontrivial nonempty types of rows and four nontrivial nonempty types of columns in any 1's cover by rectangles).

The third matrix above shows that Technique 1 is not sufficient for converting an arbitrary nonnegative rank 3 decomposition into a 1's partition with the same number of rectangles, since this matrix has a nonnegative rank decomposition with three matrices, no two of whose supports are compatible.

3.4 Separations

Proof of Theorem 2(2). Consider the following 5×6 matrix M consisting of all possible columns having a 1 in the bottom row and two 1's among the top four rows. (In fact, any column of M can

be deleted to yield a 5×5 matrix that still works, but it seems cleaner to include all the columns.)

$$M = \begin{bmatrix} 1 & 1 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & 1 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix} \qquad U = \begin{bmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

We have $\text{rank}_+(M) \geq 4$ since M has a fooling set of size 4. We also have $\text{rank}_+(M) \leq 4$ since $M = UV$ where U is as above and V is the top four rows of M. Clearly $\text{rank}_{0,1}(M) \leq 5$, so we just need to argue that $\text{rank}_{0,1}(M) \geq 5$. There are several (ad hoc) ways to see this; our preferred way is as follows. Note that the 1's of the top four rows can be partitioned into three fooling sets of size four, indicated by the colors red, green, and blue.[5] Suppose for contradiction there is a partition of the 1's of M into just four rectangles. Then each of those rectangles must contain exactly one red 1, one green 1, and one blue 1, and hence must contain all three 1's in one of the top four rows (since these are the only size-three rectangles within the top four rows). Hence each pair of rectangles in the partition shares one column in common. Thus only one of the four rectangles can touch the bottom row; but since this rectangle is three columns wide, that leaves three 1's in the bottom row uncovered, which is a contradiction. □

The upper bound $\text{rank}_+(M) \leq 4$ for Theorem 2(2) can be phrased in terms of a sampling protocol as follows. One party (it does not matter which) picks one of the top four rows uniformly at random and sends this index to the other party (so 2 bits of communication, 4 possible transcripts). Then independently of each other: Alice outputs the chosen row with probability 2/3 and the bottom row with probability 1/3. Bob outputs uniformly at random one of the three columns in which the chosen row has a 1 in M.

Proof of Theorem 2(3). Consider the following 12×12 matrix M.

$$M = \begin{bmatrix} 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 1 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 1 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 & 1 \\ 1 & 1 & 1 & 0 & 0 & 0 & 1 & 1 & 1 & 1 & 1 & 1 \end{bmatrix} \qquad U = \begin{bmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 \\ 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & \frac{1}{2} & 0 & 0 & 0 & \frac{1}{2} & \frac{1}{2} & \frac{1}{2} \end{bmatrix}$$

We have $\text{rank}_+(M) \geq 9$ since M has a fooling set of size 9. We also have $\text{rank}_+(M) \leq 9$ since $M = UV$ where U is as above and V is the top nine rows of M. Clearly $\text{rank}_{0,1}(M) \leq 12$, so we just need to argue that $\text{rank}_{0,1}(M) \geq 12$. Suppose for contradiction there is a partition of the 1's of

[5] In case the colors are not visible: red is entries $(1,1), (2,4), (3,2), (4,5)$, green is entries $(1,2), (2,1), (3,6), (4,3)$, and blue is entries $(1,3), (2,5), (3,4), (4,6)$, where row 1 is topmost and column 1 is leftmost.

M into just 11 rectangles $R_1, \ldots, R_{11}$. Let $Q_r = \{1,2,3,10\} \times \{1,2,3\}$ be the red rectangle of M, let $Q_g = \{4,5,6,11\} \times \{4,5,6\}$ be the green rectangle of M, and let $Q_b = \{7,8,9,12\} \times \{7,8,9\}$ be the blue rectangle of M. Note that these three rectangles form a fooling set. (We say $A \times B$ and $C \times D$ are fooling if either $A \times D$ or $C \times B$ is an all 0's submatrix.) Hence no R_i can intersect more than one of Q_r, Q_g, Q_b. By the pigeonhole principle, at least one of Q_r, Q_g, Q_b, say Q_r, intersects at most three of the R_i's. By inspection, the only way to partition the 1's of Q_r using at most three rectangles is to use $\{1,2,10\} \times \{1\}$, $\{1,3,10\} \times \{2\}$, and $\{2,3,10\} \times \{3\}$. Hence three of the R_i's, say R_1, R_2, R_3, intersect Q_r in those three subrectangles. Thus since each of R_1, R_2, R_3 touches two of the first three rows of M, none of R_1, R_2, R_3 can touch any of the last three columns of M (otherwise a 0-entry would be covered by an R_i). Thus the 1-entries $M_{1,10}$, $M_{2,11}$, $M_{3,12}$ must be covered by some other distinct R_i's, say R_4, R_5, R_6, none of which can intersect Q_g or Q_b since $\{1,2,3\} \times \{4,5,6,7,8,9\}$ is all 0's. This leaves $R_7, \ldots, R_{11}$ to cover the 1's of Q_g and Q_b, which is a contradiction since three rectangles are needed for Q_g, and another three are needed for Q_b. □

The upper bound $\text{rank}_+(M) \leq 9$ for Theorem 2(3) can be phrased in terms of a sampling protocol as follows. One party (it does not matter which) picks $i \in \{1, \ldots, 9\}$ uniformly at random and sends i to the other party. Then independently of each other: Alice outputs row i with probability 1/2, and with probability 1/4 each outputs one of the two rows in which column i has a $\frac{1}{2}$ in U. Bob outputs uniformly at random one of the three columns in which row i has a 1 in M.

4 Open Questions

Can Theorem 1 or Lemma 2 be quantitatively improved? Can it be witnessed by a composed function with a simpler gadget than the confounding gadget (and with an elementary proof that avoids the machinery of [GLM+15, Göö15])?

What is the best separation between log of nonnegative rank and log of binary rank for total 0-1 matrices? We know it is at least an additive constant and at most quadratic.

A number of so-called simulation theorems are known, which convert lower bounds for query complexity measures into lower bounds for the corresponding communication complexity measures (e.g., [She11, GLM+15, GPW15]). Unfortunately, no such simulation theorem is known for unambiguous nondeterminism; such a result could be useful for obtaining new binary rank lower bounds.

For total 0-1 matrices M, N, let $M \otimes N$ denote the (Kronecker) tensor product, and let $M^{\otimes k}$ denote the k-fold tensor product of M with itself. How do rank_+ and $\text{rank}_{0,1}$ behave under the tensor product? This can be viewed as a direct sum question. We trivially have

$$\max\big\{\text{rank}_+(M) \cdot \text{fool}(N),\ \text{fool}(M) \cdot \text{rank}_+(N)\big\} \ \leq\ \text{rank}_+(M \otimes N) \ \leq\ \text{rank}_+(M) \cdot \text{rank}_+(N)$$

and

$$\max\big\{\text{rank}_{0,1}(M) \cdot \text{fool}(N),\ \text{fool}(M) \cdot \text{rank}_{0,1}(N)\big\} \ \leq\ \text{rank}_{0,1}(M \otimes N) \ \leq\ \text{rank}_{0,1}(M) \cdot \text{rank}_{0,1}(N)$$

where $\text{fool}(M)$ denotes the largest size of a fooling set of M. It does not seem to be known whether there exists an M and k for which $\text{rank}_{0,1}(M^{\otimes k}) < \text{rank}_{0,1}(M)^k$. It is known [BKLT13] that there exist total real nonnegative matrices M, N such that $\text{rank}_+(M \otimes N) < \text{rank}_+(M) \cdot \text{rank}_+(N)$.

Conjecture 1. *For all k, $\text{rank}_{0,1}(M^{\otimes k}) = 5^k$ where matrix M is from the proof of Theorem 2(2).*

For this particular M, we have $\text{rank}_+(M^{\otimes k}) = \text{fool}(M^{\otimes k}) = 4^k$ for all $k \geq 1$. Hence under this conjecture, we would have a family of total matrices for which unambiguous nondeterminism requires $\log_4(5) \approx 1.16$ times more bits of communication than the corresponding sampling problem. We have been unable to make any significant progress on the conjecture.

Is there some other way (besides possibly tensor product) to amplify the gap and get a family of total 0-1 matrices exhibiting some separation between nonnegative rank and binary rank? It is conceivable that a computer search could be used to find a better gap example (e.g., by confirming Conjecture 1 for $k = 2$). However, we feel that such an enterprise would be unenlightening and unlikely to lead to any proof techniques for obtaining a superlinear gap between nonnegative rank and binary rank.

A measure sandwiched between rank_+ and $\text{rank}_{0,1}$ is the number of rectangles needed to uniformly cover the 1's of a matrix (so all 1's are covered by the same number of rectangles). How does this measure relate to rank_+ and $\text{rank}_{0,1}$? The upper bound in Theorem 1 also works for this intermediate measure, but the upper bounds in Theorem 2(2) and Theorem 2(3) do not.

Acknowledgments

I thank Mika Göös, Troy Lee, and Toniann Pitassi for discussions and anonymous reviewers for comments.

References

[AKK16] Andris Ambainis, Martins Kokainis, and Robin Kothari. Nearly optimal separations between communication (or query) complexity and partitions. In *Proceedings of the 31st Computational Complexity Conference (CCC)*, 2016. To appear.

[ASTS+03] Andris Ambainis, Leonard Schulman, Amnon Ta-Shma, Umesh Vazirani, and Avi Wigderson. The quantum communication complexity of sampling. *SIAM Journal on Computing*, 32(6):1570–1585, 2003.

[BKLT13] LeRoy Beasley, Hartmut Klauck, Troy Lee, and Dirk Oliver Theis. Communication complexity, linear optimization, and lower bounds for the nonnegative rank of matrices (Dagstuhl seminar 13082). *Dagstuhl Reports*, 3(2):127–143, 2013.

[CR93] Joel Cohen and Uriel Rothblum. Nonnegative ranks, decompositions, and factorizations of nonnegative matrices. *Linear Algebra and Its Applications*, 190:149–168, 1993.

[Gil14] Nicolas Gillis. The why and how of nonnegative matrix factorization. In *Regularization, Optimization, Kernels, and Support Vector Machines*, pages 257–291. Chapman & Hall/CRC, 2014.

[GLM+15] Mika Göös, Shachar Lovett, Raghu Meka, Thomas Watson, and David Zuckerman. Rectangles are nonnegative juntas. In *Proceedings of the 47th ACM Symposium on Theory of Computing (STOC)*, pages 257–266, 2015.

[Göö15] Mika Göös. Lower bounds for clique vs. independent set. In *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1066–1076, 2015.

[GPW15] Mika Göös, Toniann Pitassi, and Thomas Watson. Deterministic communication vs. partition number. In *Proceedings of the 56th IEEE Symposium on Foundations of Computer Science (FOCS)*, pages 1077–1088, 2015.

[JSWZ13] Rahul Jain, Yaoyun Shi, Zhaohui Wei, and Shengyu Zhang. Efficient protocols for generating bipartite classical distributions and quantum states. *IEEE Transactions on Information Theory*, 59(8):5171–5178, 2013.

[Juk12] Stasys Jukna. *Boolean Function Complexity – Advances and Frontiers*, volume 27 of *Algorithms and Combinatorics*. Springer, 2012.

[KLO99] Eyal Kushilevitz, Nathan Linial, and Rafail Ostrovsky. The linear-array conjecture in communication complexity is false. *Combinatorica*, 19(2):241–254, 1999.

[KN97] Eyal Kushilevitz and Noam Nisan. *Communication Complexity*. Cambridge University Press, 1997.

[KW09] Eyal Kushilevitz and Enav Weinreb. On the complexity of communication complexity. In *Proceedings of the 41st ACM Symposium on Theory of Computing (STOC)*, pages 465–474, 2009.

[Lov14] Shachar Lovett. Communication is bounded by root of rank. In *Proceedings of the 46th ACM Symposium on Theory of Computing (STOC)*, pages 842–846, 2014.

[LS93] László Lovász and Michael Saks. Communication complexity and combinatorial lattice theory. *Journal of Computer and System Sciences*, 47(2):322–349, 1993.

[NW95] Noam Nisan and Avi Wigderson. On rank vs. communication complexity. *Combinatorica*, 15(4):557–565, 1995.

[SA15] Manami Shigeta and Kazuyuki Amano. Ordered biclique partitions and communication complexity problems. *Discrete Applied Mathematics*, 184:248–252, 2015.

[She11] Alexander Sherstov. The pattern matrix method. *SIAM Journal on Computing*, 40(6):1969–2000, 2011.

[Shi14] Yaroslav Shitov. An upper bound for nonnegative rank. *Journal of Combinatorial Theory, Series A*, 122:126–132, 2014.

[Vio12] Emanuele Viola. The complexity of distributions. *SIAM Journal on Computing*, 41(1):191–218, 2012.

[Yan91] Mihalis Yannakakis. Expressing combinatorial optimization problems by linear programs. *Journal of Computer and System Sciences*, 43(3):441–466, 1991.